Jureca

JÜLICH FORSCHUNGSZENTRUM

Computational Ecosystem in Neuroscience:

High Performance Analytics and Computing for the HBP

3 Nov 2018 | Alexander Peyser (a.peyser@fz-juelich.de) | Simulation Lab Neuroscience
Jülich Supercomputing Centre
Institute for Advanced Simulation
Forschungszentrum Jülich

The Human Brain Project

- HBP is a Future and Emerging Technologies (FET) Flagship project co-funded by the European Union with a total duration of 10 years and up to 1 billion € funding
- $>$ 100 partner institutions from 20 European countries (and others)

Goals for 10 year endeavor

- Build a European Research Infrastructure (RI) for the neuroscience community, including six platforms:
  Neuroinformatics, Brain Simulation, High Performance Analytics and Computing, Medical Informatics, Neuromorphic Computing, Neurorobotics
- Gather, organize and disseminate data describing brain and its diseases
- Construct large scale brain simulation
- Build multi-scale scaffold theory and models of the brain
- Develop brain-inspired computing, data analytics and robotics.

High Performance Computing and Analytic Platform

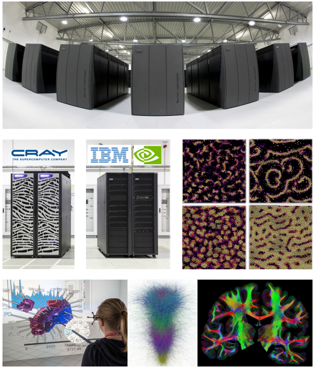

We build and operate a supercomputing, data and visualization infrastructure enabling scientists to

- Run large-scale, data intensive, interactive brain simulations up to the size of a full human brain
- Manage the large amounts of data used and produced in the HBP
- Manage complex workflows comprising concurrent simulation, data analysis and visualization workloads

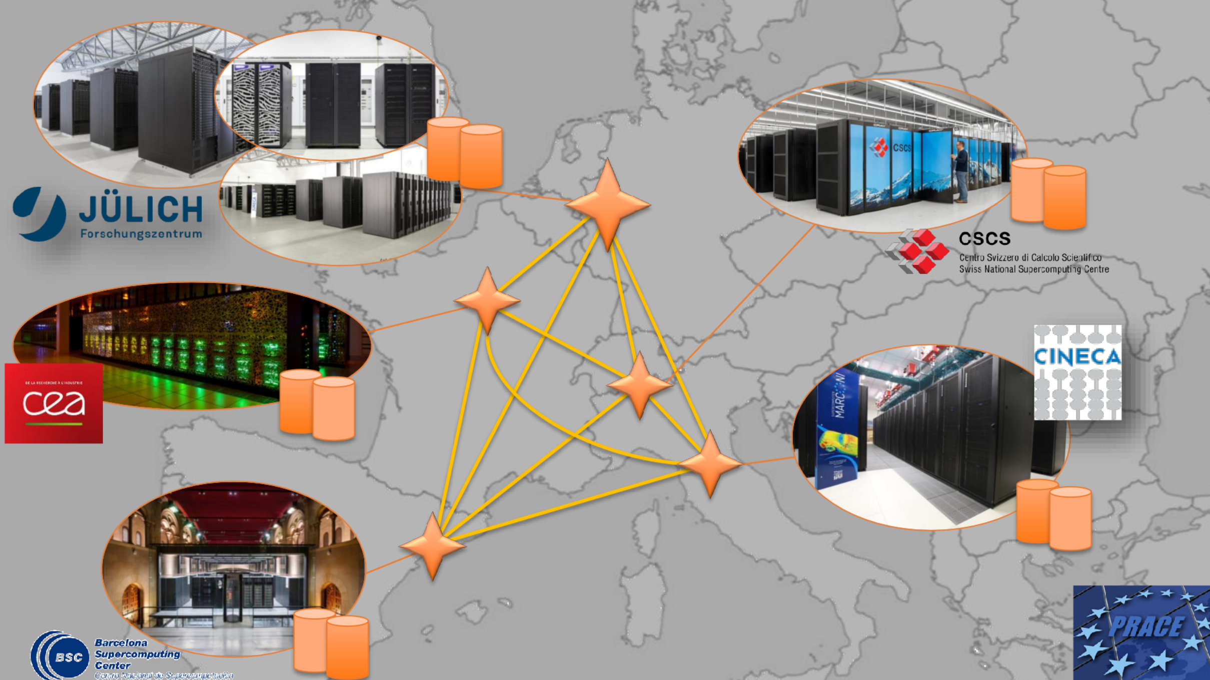

JÜLICH
Forschungszentrum
CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre
CINECA
cea
Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación
PRACE

Fenix

Federated data & compute infrastructure

Aim	Providing scalable compute and data services in a federated manner
Target communities	• Neuroscience: HBP is prioritized driver for infrastructure design • Materials science • Genomics • Physical science experiments • … and others with similar requirements
Key features	• Scalable compute resources • Federated data infrastructure • Interactive Compute Services providing access to the federated data infrastructure as well as elastic access to the scalable compute resources
Consortium	

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

cea

CINECA

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

JÜLICH
FORSCHUNGSZENTRUM

Software stacks

FORSCHUNGSZENTRUM JÜLICH

Inputs needing analysis:

- Electrophysiological data analysis
- fMRI analysis pipelines
- EEG data
- Microscopic data for brain atlases

Simulations for model testing:

- Neural mass models: 700 frequency coupled regions for a full brain
- Point neuron models: up to 10^{11} neurons approximated as a network of time-delayed binary signals and an integration function defining firing times
- Morphologically detailed neuron models: neurons are modeled as combinations of synaptic functions, coupled passive cables with active variable resistors
- ... and at the bottom, reaction-diffusion models for synapses, signal exchanges and ion channels
- Multiscale models combining fast, coarse-grain models for global behavior and fine-grain models for perturbation-sensitive regions

Large scale analysis & machine learning

- Workflow construction toolset: Modular Science
- Parameter optimization tools: Learning-2-Learn, JuPeX

Visualization:

- Analytic pipelines visualization
- Analysis for network construction
- Interactive and steered simulations
- Experiment & provenance tracking

Human Brain Atlas

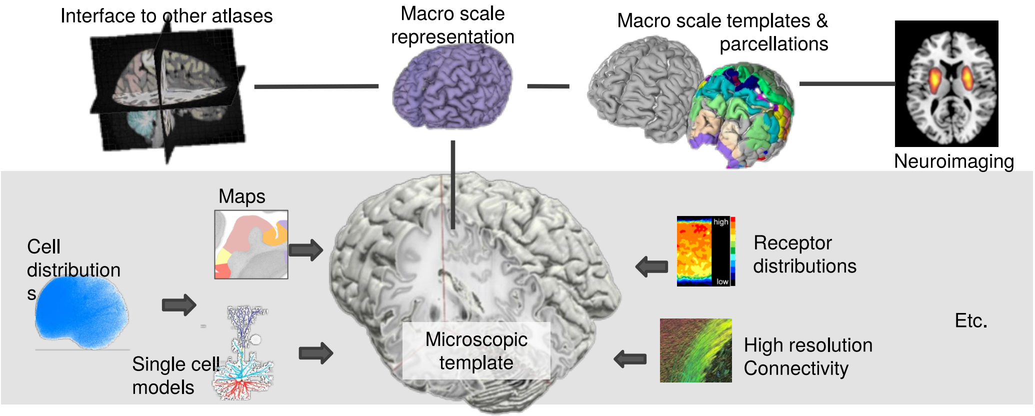

Timo Dickscheid

External data to analysis

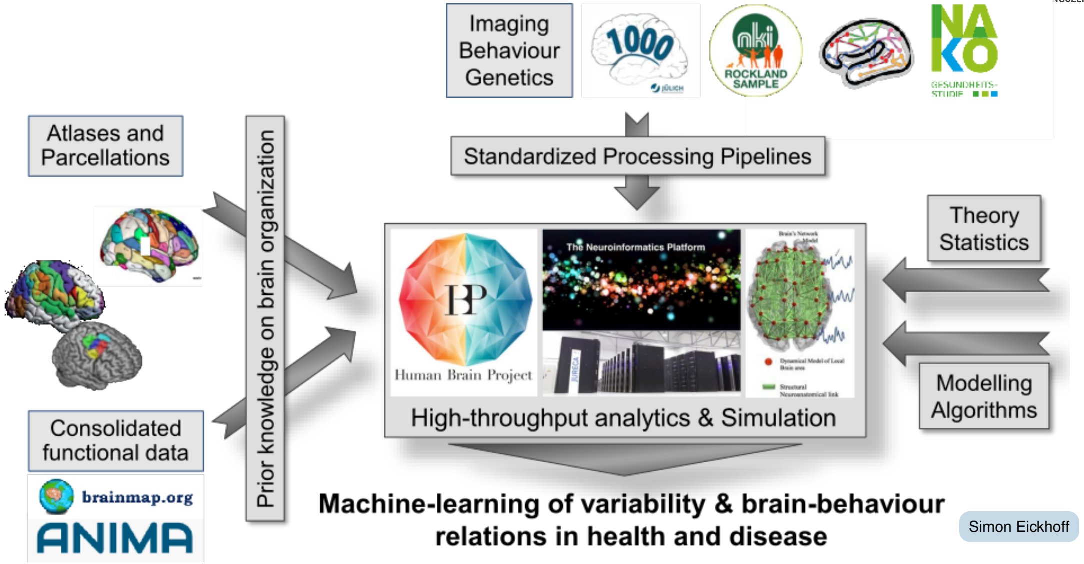

The Virtual Brain

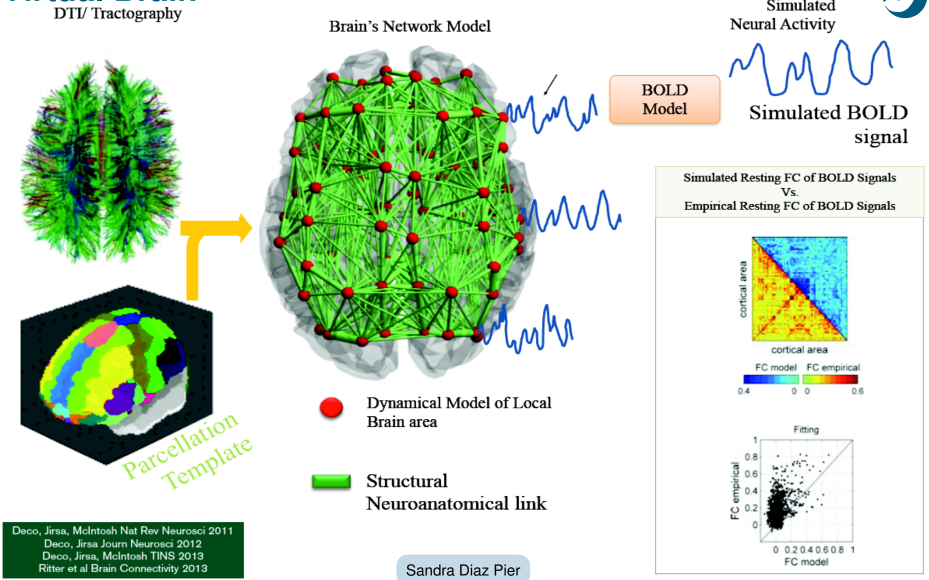

Deco, Jirsa, McIntosh Nat Rev Neurosci 2011
Deco, Jirsa Journ Neurosci 2012
Deco, Jirsa, McIntosh TINS 2013
Ritter et al Brain Connectivity 2013

Sandra Diaz Pier

TVB for HPC

"The Virtual Brain" software package for solving these systems is a combination of Matlab and Python scripts attached to a front-end visual control center, all which are designed for tradition desktops. To do parameter searches over individual dynamics for large data sets requires HPC resources

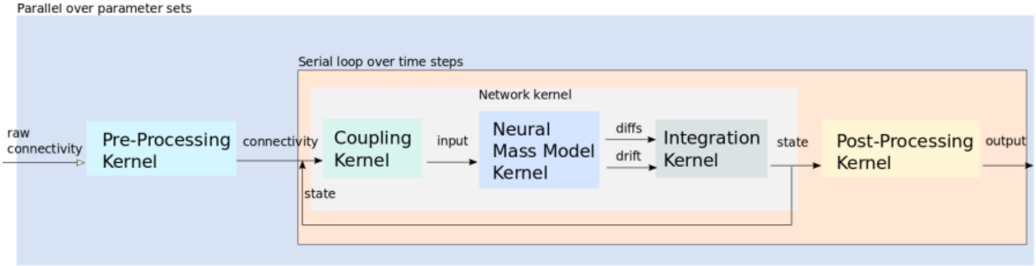

Approach:

Sandra Diaz Pier, Marmaduke Woodman (Blundell et al 2018)

- Automatic kernel code generation using 'loo.py'
- Domain Specific Language to describe neural mass models, coupling, integration, and pre- and post-processing kernels.
- Use one description for multiple backends: CPUs, GPUs, KNL, FPGAs, ...

Collaboration w/ Institut des Neuroscience des Systèmes, Aix-Marseille Université

NEST = NEural Simulation Tool

Simulations of large networks of spiking point-neuron models (internal single-neuron dynamics can be arbitrarily complex)

- C++ kernel
- built-in simulation language interpreter (SLI)
- Python-based user interface (PyNEST)
- NESTml neuron modeling DSL
- multi-threading to use multi-processor machines efficiently
- MPI & OpenMP parallelism to use compute clusters and supercomputers

- Synaptic plasticity
- Structural plasticity
- Electrically coupled neurons

arbor: Background

Arbor is a project to develop:

- a new morphologically detailed, multi-compartmental neuronal network simulator library,
- that is optimized for HPC systems,
- and is easy to integrate into existing workflows.

Developed

- between the Jülich Supercomputing Centre at the Forschungszentrum Jülich and CSCS the Swiss Supercomputing Centre as a component of the Human Brain Project
- as an open-source project
- using an open developmental model

See current development at

- `https://github.com/arbor-sim/arbor`

arbor: Performance

daint-mc	Cray XC40: 2× 18-core Broadwell per node
daint-gpu	Cray XC50: 1× P100 GPU per node
cells	150 compartments & 10,000 synapses per cell Passive dendrites, Hodgkin-Huxley soma
network	ring network
duration	100 ms

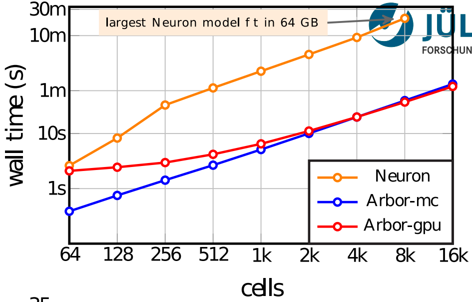

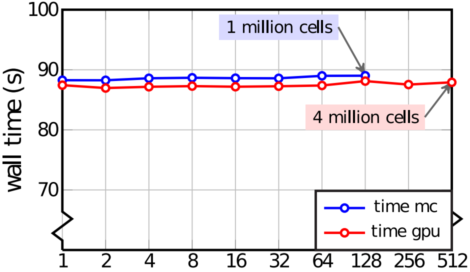

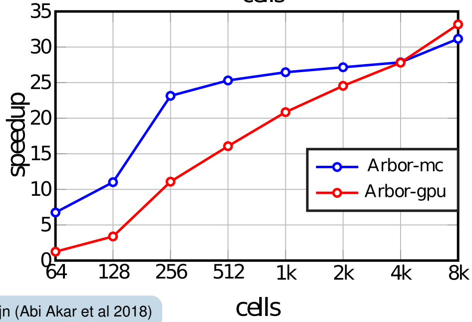

Ben Cumming, Sam Yates, Nora Abi Akar, Anne Küsters, Wouter Klijn (Abi Akar et al 2018)

arbor: GPU optimization

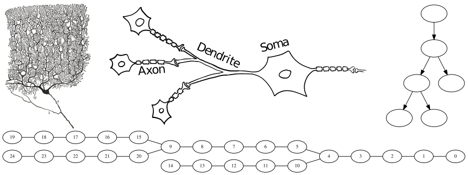

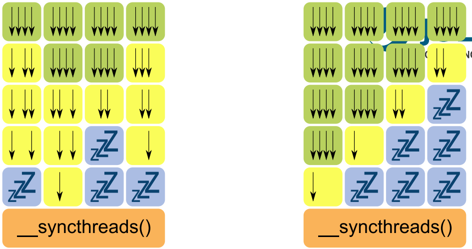

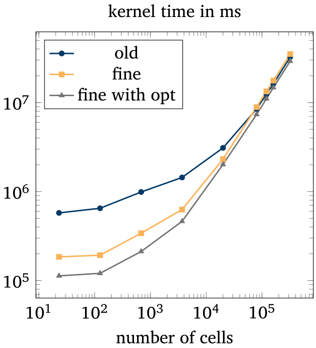

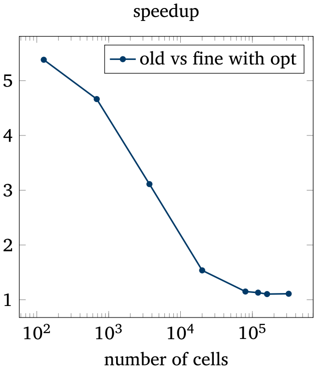

Felix Huber 2018

Modular Workflows

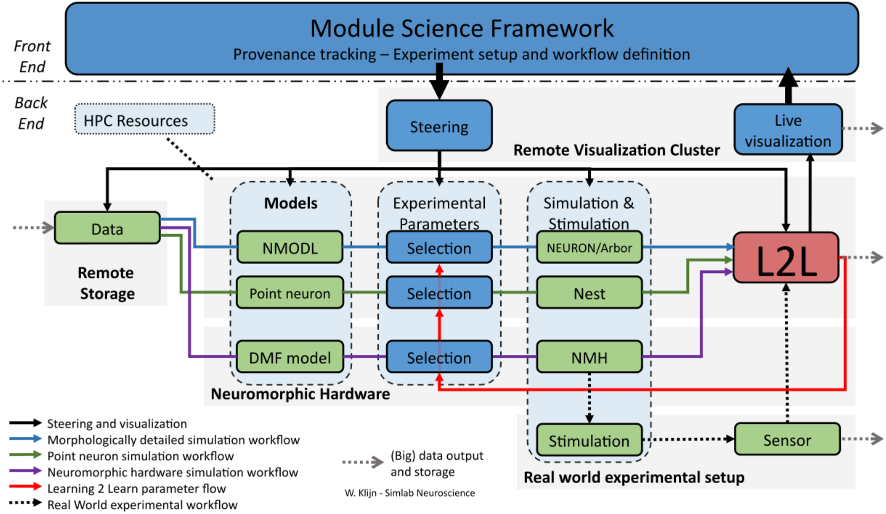

Wouter Klijn

Visualization Framework

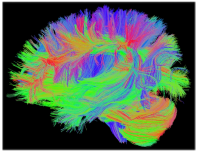

Caspers *et al*. 2014

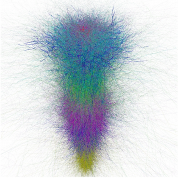

Pastor *et. al.* 2015, Data: BBP

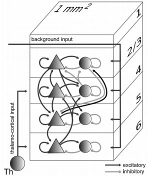

Potjans *et. al.* 2011

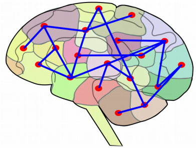

Woodman, *et al.* 2017

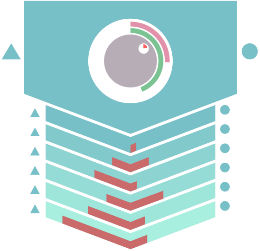

NeuroScheme

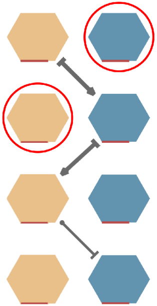

Learning2Learn Concept

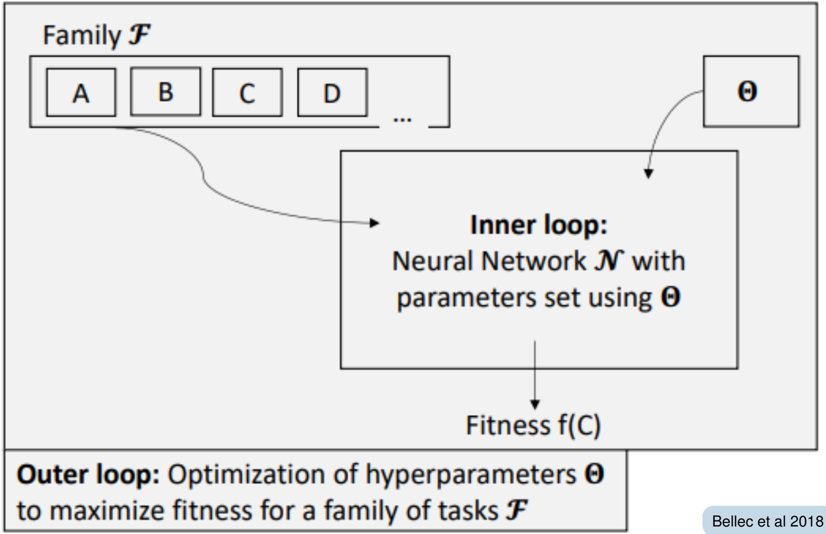

Bellec et al 2018

Learning2Learn Architecture

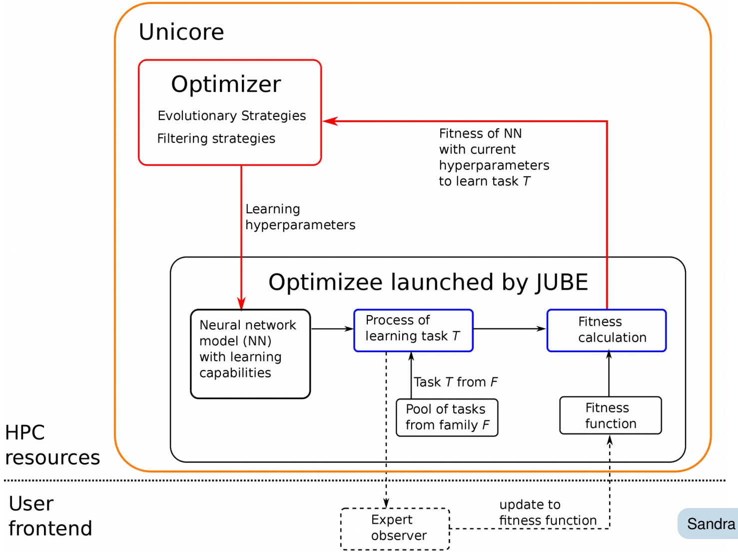

Sandra Diaz Pier

Workshops

Symposium on generative connectomics and plasticity

Date: March 26–29 2019
Location: Jülich Germany
Join us in looking at the state of the art in simulation and modeling of generative connectomics and plasticity. We will discuss modeling and simulation of connectivity generation from two perspectives:

- neural development and structural plasticity in biological neural networks
- generation of connectivity in spiking and artificial neural networks

HPAC training

Date: December 11 & 12
Location: Barcelona
Training in HPAC infrastructure tools and a quick introduction to NEST & Arbor

Proposed: High Performance Computing for Neuroscience: Hands-on introduction to supercomputing usage, tools and applications

Date: June ?? (3 day course)
Location: Jülich Germany
This workshop will cover the basics of supercomputing needed for getting started using HPC systems for (neuroscience) research. Including: Python, machine usage, data management, introduction to NEST, Arbor and visualization, and usage of other HBP platform tools such as neurorobotics software and neuromorphic hardware in the context of HPC.

JÜLICH
FORSCHUNGSZENTRUM

We're hiring

Are you a computer scientist or mathematician interested in all aspects of computational neuroscience at the intersection of high-performing computing?

Can you solve inverse problems, program GPUs or love to build systems with exotic neuromorphic hardware or build complex neuroscientific multiscale workflows?

The SimLab Neuroscience at the Jülich Supercomputing Center in the Forschungszentrum Jülich is hiring multiple positions.

Please contact a.peyser@fz-juelich.de.

Thank you!

JÜLICH FORSCHUNGSZENTRUM

And thanks to...

SimLab

Sandra Diaz Piers
Wouter Klijn
Guido Trensch
Jochen Eppler
Lekshmi Deepu
Anna Lührs
Meredith Peyser
Abigail Morrison
Boris Orth
Bastian Tweddell
Kai Krajsek
Anne Küsters
Anne Carstens
Andreas Herten

...and everyone else in the SimLab and HPCNS

Arbor: CSCS

Ben Cumming
Sam Yates
Nora Abi Akar
Marco Kabic
Vasileos Karakakis

TVB external

Viktor Jirsa, Marseille
Marmaduke Woodhouse, Marseille
Petra Ritter, Charité
Michael Schirner, Charité
Olaf Sporns, Indiana University

... and everyone else across the HBP

Imaging: Institute for Neuroscience & Medicine

Katrin Amunts
Timo Dickscheid
Simon Eickhoff
Thanos Manos
Felix Hoffstaedter
Popovych, Oleksandr

NEST: NMBU

Susanne Kunkel
Hans Ekkehard Plesser

Graz Technische Universität

Wolfgang Maas
Anand Subramoney

References

N. Abi Akar, J. Biddiscombe, B. Cumming, M. Kabic, V. Karakasis, W. Klijn, A. Küsters, I. Martinez, A. Peyser, and S. Yates. arbor-sim/arbor: Version 0.1: First release; v0.1, 2018. doi: 10.5281/zenodo.1459679.

G. Bellec, D. Salaj, A. Subramoney, R. Legenstein, and W. Maass. Long short-term memory and learning-to-learn in networks of spiking neurons. 2018. URL https://arxiv.org/abs/1803.09574.

I. Blundell, R. Brette, T. A. Cleland, T. G. Close, D. Coca, A. P. Davison, S. Diaz-Pier, C. Fernandez Musoles, P. Gleeson, D. F. M. Goodman, M. Hines, M. W. Hopkins, P. Kumbhar, D. R. Lester, B. Marin, A. Morrison, E. Müller, T. Nowotny, A. Peyser, D. Plotnikov, P. Richmond, A. Rowley, B. Rumpe, M. Stimberg, A. B. Stokes, A. Tomkins, G. Trensch, M. Woodman, and J. M. Eppler. Code generation in computational neuroscience: A review of tools and techniques. Frontiers in Neuroinformatics, 12:68, 2018. ISSN 1662-5196. doi: 10.3389/fninf.2018.00068.

F. Huber. Efficient tree solver for hines matrices on the gpu using fine grained parallelization and basic work balancing. In Guest Student Programme on Scientific Computing. Forschungszentrum Jülich, 2018. URL https://arxiv.org/abs/1810.12742.

daint-mc	Cray XC40: 2 × 18-core Broadwell per node
daint-gpu	Cray XC50: 1 × P100 GPU per node
cells	150 compartments & 10,000 synapses per cell Passive dendrites, Hodgkin-Huxley soma
network	ring network
duration	100 ms